

## MODIS products and MODAPS computing systems

This short write-up presents information about MODIS the volume of MODIS products and processing rates achieved in the MODAPS system. It also provides a brief overview of the individual computing systems that make up MODAPS and how these are maintained. It is accurate as of February 2008.

### MODIS Product Volumes

Average daily data volumes for MODIS products produced at 1x (1 MODIS instrument 1 day's worth of products) are:

Level 0 Raw data 71GB

Level 1 (Calibrated Radiances and Geolocation) 210 GB

Level 2 (Geophysical parameters in instrument swath format) Atmosphere products 10 GB

Level 2 Land products 70 GB

Level 3 Atmosphere products (global grid) 0.5 GB

Level 3 Land products (global grid) 196 GB

### MODIS Processing and Distribution Rates

MODAPS system currently process at rates of 2x for current forward processing (all products for both MODIS instruments) with additional catch-up capability to avoid falling behind. In addition to forward processing performed on one of the three processing strings, two additional production strings are used for reprocessing of the MODIS data record (Terra MODIS Feb 2000 to present, Aqua MODIS July 2002 to present) at rates of 25x for Atmosphere products and 12x for Land products and for large-scale science testing. Additional compute servers to be installed in the 2008 are expected to increase the reprocessing rate to approximately 40x for Atmosphere products and 20x for Land products.

After MODIS products are produced by the MODAPS system, the products are sent to the DAACs and moved into the LAADS system for archiving and subsequent distribution to the public. The corresponding volumes for 1x are presented below. Product volumes denoted as COMBINED are for those products that use input from both MODIS instruments to make a single product.

Exported to Land Processes DAAC:

L2: 1.1 GB/DataDay per satellite  
L3: 57.1 GB/Dataday per satellite  
6.1 GB/Dataday for COMBINED

Exported to NSIDC:

L2: 0.6 GB/Dataday per satellite  
L3: 0.4 GB/Dataday per satellite

Archived to LAADS:

L2 Atmos: 10.3 GB/Dataday per satellite  
L3 Atmos: 0.6 GB/Dataday per satellite  
  
L2 Land: 0.4 GB/Dataday per satellite  
L3 Land: 20.1 GB/DataDay per satellite  
7.0 GB/DataDay for COMBINED

#### Links to information about MODIS products

More information about specific MODIS products can be obtained from the Land, Atmosphere and Calibration sites linked to the MODIS home page: <http://modis.gsfc.nasa.gov/> and from the Land Processes DAAC <http://edcdaac.usgs.gov/index.asp>, the National Snow and Ice Data Center DAAC <http://nsidc.org/daac/modis/index.html> and the Level 1 and Atmosphere Archive and Distribution System (LAADS) site <http://ladsweb.nascom.nasa.gov/>. Please note that the MODIS Ocean products are produced by the Ocean Color Data Processing System (OCDPS) whose work is not included under this contract.

#### Overview of MODAPS computing systems

Information about the MODAPS Linux computing systems and their peripherals appears in Table 1. A short description of the categories in that list is as follows:

- Minions – Linux processing servers used to produce products

- Level 0 – storage servers with RAID disk to hold Level 0 (raw) MODIS products
- Datapool – Linux servers and storage used to hold those MODIS Level 1 calibrated radiance products that are not permanently archived, for a period of 30 days.
- AADS storage nodes – Linux servers holding the archive of MODIS Atmosphere products
- LADS storage nodes – Linux servers holding the archive of MODIS Land products
- POD servers – Processing On Demand to produce products not stored in the archive and to transform products in the archive to user specified formats, projects and subsets.

Function	Quantity	Manufacturer & Model	CPU type/#	Operating System	Memory/Disk
Compute Server	77	Dell PowerEdge 1550	Intel PIII 1.1Ghz, 2	Mandrakelinux 10.1	2GB/0.26TB
Compute Server	26	Dell PowerEdge 2650	Intel Xeon 2.4Ghz, 2	Mandrakelinux 10.1	2GB/0.26TB
Compute Server	101	Dell PowerEdge 2850	Intel Xeon 2.8Ghz, 2	Mandrakelinux 10.1	2GB/0.26TB
Compute Server	100	Dell PowerEdge 2950	Intel Xeon 2.3Ghz, 2	Mandrakelinux 10.1	3.2G/0.32TB
Computer Server	100	Dell PowerEdge 2950	Intel Xeon 2.7Ghz, 2	Mandrakelinux 10.1	3.2G/0.8TB
L0 archive Server	16	Dell PowerEdge 2850	Intel Xeon 2.8Ghz, 2	Mandrakelinux 10.1	2GB/21.8TB
Atmosphere storage servers	7	Dell PowerEdge 2650	Intel Xeon 2.4Ghz, 2	Mandrakelinux 10.1	2GB/13.6-26.7TB
	6	Dell PowerEdge 2850	Intel Xeon 2.8Ghz, 2		
Land storage servers	6	Dell PowerEdge 2850	Intel Xeon 2.8Ghz, 2	Mandriva 2006	2GB/8.7-27.3TB
Datapool disk servers	11	Dell PowerEdge 1550, 2650, 2850, 6650	Intel Xeon	Mandriva 2006	2.0-3.8GB/5.5-28.3TB
Processing On-Demand Server	5	Dell PowerEdge 2850	Intel Xeon 2.8Ghz, 2	Mandriva 2006 Solaris	2GB/9.7-18.4TB
LAADS Database	1	Dell PowerEdge 6850	Intel Xeon 3.2Ghz, 8	Mandriva 2007	32GB
Production Databases	5	Dell PowerEdge 6850 Dell PowerEdge 2850	Intel Xeon 2.7Ghz, 4 Intel Xeon 2.8Ghz, 2	Mandriva 2007	16GB 8GB
Central hosts	3	Dell R900	Intel Xeon 2.1Ghz, 2	Solaris	64GB/.6TB
Additional servers – Q/A, development	7	Dell 2650, Dell 2850, Dell 6850	Intel Xeon	Mandriva	varies

Table 1. Linux computing systems in MODAPS

Assume for the purposes of sizing the staff needed to support these processing systems that an additional 140 new servers will be added in 2008 and from 200 to 500 servers will be added over the period 2009 - 2011. Additional data storage will also be added to the configuration in over this period. Assume for planning purposes that the amount of new disk storage added each year will be between 200TB and 800TB.

In addition to the Linux servers used for product generation and product distribution, there are four Silicon Graphics Incorporated Origin servers. An Origin 3800 with 64 processors and 51GB of memory and 2 Origin 300 servers each with 28 processors with 28GB of memory are used as central processing hosts. These hosts schedule jobs, stage data products for processing, run science software that has yet to be ported to Linux and push data products to science team members and other users with data subscriptions via ftp. Each Origin server has approximately 100TB of Fibre Channel attached RAID disk used to hold products from processing and large scale science tests. In addition to the 3 central processing hosts there are two smaller SGI Origin 3200 systems that are used for integration and testing by the MODAPS developers and the MODIS Science Team and an additional SGI Origin 3800 (32 processors with 32GB of memory) that is attached to a central processing host as a resource to run production jobs performing the same function as the Linux minion compute servers above. The SGI hosts are running version 6.5.22m of the Irix operating system. We anticipate adding 3-5 additional SGI Origin 3000 hosts that are being retired by other NASA projects in 2009-2010. These systems will be used as spares for existing central processing hosts and as Irix-based minions. Over time we will eliminate the Irix servers as the central processing hosts and dispatch jobs from Linux servers to a mix of Linux and Irix-based compute servers.

System administrators provide coverage for the MODAPS processing systems from 8am-10pm every day (including weekends and holidays.) For the portion of the weekdays outside of normal business hours, 8am-6pm, and on weekends, one system administrator is assigned pager duty. The operations team and software monitoring critical components of the system will send alerts to the system administrator on duty who will resolve the problem remotely if possible or come in to GSFC to deal with the problem in person if required. The components of the system are for the most part highly redundant so the loss of an individual processing server, Irix or Linux minion, is not generally addressed until the next business day. Critical failures that impact:

1. overall processing on a production string, such as the processing database servers for each central production host, the central host itself or a critical piece of the network infrastructure; or
2. distribution of products, such as the distribution database server, web server or ftp server, network components or the central ftp servers for distribution of products to the EOS DAACs

are addressed as soon as a page is received. Critical components are typically restored within 2 hours. However, full restoration of functionality may take up to several days when spares are not stocked onsite.

There are no service contracts on the SGI Origin systems and all maintenance of both software and hardware for these systems is handled by the onsite system administrators. There is warranty coverage for Dell servers purchased within the last 3 years (models 2850, 2950 and 6850 in Appendix C) and system administrators contact the vendors for replacement parts that are cross shipped. All other Dell servers are out of warranty and system administrators repair failed systems by swapping parts from other failed systems until all components in a system are bad at which time it is retired.

Warranty coverage (5 years) exists for: Seagate Fibre Channel, SATA and SAS disk drives that we have purchased and for RAID Inc Falcon storage arrays. System administrators handle cross shipping of replacement parts with storage vendors. Data Direct Networks (DDN) disk arrays are used in the production of MODIS products on the central processing hosts and in archiving the Level 0 raw data products. The earliest versions of DDN disk arrays, S2A6000 controllers, enclosures and disks, are not under warranty. All maintenance of these storage systems is handled by the system administration team using onsite spares and disk drives purchased and installed in the arrays (Seagate Fibre Channel drives mentioned in previous paragraph). Newer models of DDN disk arrays, the S2A8000 and S2A9000, are under maintenance with DDN which ships spare parts to our facility on request as well as provides firmware updates. Onsite maintenance of these storage arrays is handled by members of the system administration team who have received DDN training and/or in-house training by the MODAPS system engineer.

Maintenance is in place for the three enterprise class Extreme Networks Black Diamond switches (models BD 6816 and BD 10808) that handle the main traffic for MODIS production and distribution. There is also an older Extreme Networks Alpine class switch that handles traffic for approximately 100 Dell servers (Dell 1550 servers in Appendix A) that is connect to the central processing hosts. The remaining 40 smaller switches (24 ports or less) from Extreme Networks, Cisco, Brocade, Hewlett Packard and Linksys are covered under a combination of warranties and shelf spares.

The tape units that are currently maintained by the system administration staff include:

1. 1 ADIC Scalar 1000 tape library used by Land Q/A staff to backup database entries
2. A pair of LTO 2 tape drives used for limited product distribution to scientists

Additional systems and peripherals in facilities supported by system administrators include:

1. Backup server in Building 33
2. 2 network firewalls, 1 mail server and 25 Linux workstations at an offsite facility
3. Account management server
4. Color and black and white laser printer
5. Large format color inkjet plotter, HP DesignJet 3500CP